

CONTRIBUTION TO T1 STANDARDS PROJECT

TITLE: A proposal for SONET Standards on Virtual Concatenation of High Order and Low Order SPEs

STANDARDS PROJECT: Digital Optical Hierarchy

PD.
P: 1.12 (12)

AUTHOR: Nevin Jones (Lucent Technologies) and Michael Mayer (Nortel Networks)

SOURCE:

Lucent Technologies
Bell Labs Innovations



NORTEL
NETWORKS

CONTACT:

S. J. (John) Chen
Lucent Technologies
Room 4L-322
101 Crawfords Corner Road
Holmdel, NJ 07733 USA
Tel: 732 949 5179
Fax: 732 949 5055
Email: sichen@lucent.com

Michael Mayer
Nortel Networks
P.O. Box 402
Ogdensburg NY 13669 USA
Tel: 613-765-4153
Fax: 613-763-2388
E-mail: mgm@nortelnetworks.com

DATE: January 17th-21st, 2000

DISTRIBUTION: T1X1.5

ABSTRACT

In the interest of providing a common global standard, this contribution proposes that committee T1X1 begin work on the standardization of virtual concatenation of SONET STS-n and VT-n SPEs in alignment with the work that has been successfully done by SG-15/Q9 and SG-15/Q11 of the ITU-T.

T1X1.5/99-XXX

NOTICE

The proposals in this submission have been formulated to assist ATIS subcommittee T1X1. This document is offered to the subcommittee as a basis for discussion and is not a binding proposal on either Lucent Technologies or Nortel Networks. The requirements are subject to change in form and numerical value after more study. Lucent Technologies and Nortel Networks specifically reserve the right to add to, or amend the statements made herein. Nothing contained herein shall be construed as conferring by implication, estoppel or otherwise any license or right under any patent, whether or not the use of any information herein necessarily employs an invention of any existing or later issued patent.

BEST AVAILABLE COPY

Introduction

It is now a common place notion that data traffic on the public network is growing at exponential rates and is forecast to surpass traditional voice traffic in the 2002-2005 time-frame.

Network operators such as France Telecom, Deutsche Telecom and British Telecom are planning an early roll-out of equipment supporting virtual concatenation to provide for optimal transport of these emerging data services.

It is therefore not accidental that such operators have been very active in fora that have focused on virtual concatenation standardization efforts. This is especially true with respect to the ETSI TM1/WP3 meeting (August 31 to September 3, 1999 at Nice) and the ITU-T SG-15/Q9 and SG-15/Q11 Rapporteur meeting (November 15-19, 1999 at Ipswich) where this technique achieved its most significant progress through the standardization process.

Discussion

The ETSI TM1/WP3 meeting resulted in a virtually unanimous support for standardization work on virtual concatenation and liaised several statements of support (to the ITU-T SG-15 and SG-13 bodies) for standardization of virtual concatenation for High Order (HO) and Low Order (LO) SDH payloads.

At a November 15-19, 1999 ITU-T SG-15/Q11/Q9 Rapporteur meeting held at at Ipswich UK, agreement was reached on supporting the development of standards for HO and LO virtual concatenation.

The baseline text with the technical details for HO virtual concatenation is contained in the current drafts of G.707 and G.783. The final text is scheduled for inclusion in the March 2000 update of G.707 (see Appendix A), G.783 and other related ITU-T standards.

Final text for the technical details for LO virtual concatenation, is scheduled to be available by the end of January 2000 for inclusion into the March 2000 update of G.707, G.783 and other related ITU-T standards.

Proposal

Over the last several months, a number of contributions (e.g., [1] and [2]) detailing the benefits of SONET HO and LO virtual concatenation have been made before Committee T1X1.

This contribution submits that the next step must be the establishment of a work effort to develop SONET standards for virtual concatenation.

The following proposal is therefore made:

1. An identical methodology as that currently contained in G.707 and G.783 to implement virtual concatenation of VC-3/VC-4 payloads should be developed in SONET for STS-n SPEs.
2. An identical methodology as that to be included in the March 2000 update of G.707 and G.783 to implement virtual concatenation of VC-11/VC-12 payloads, should be developed in SONET for VT-n SPEs.
3. T1X1 should develop and bring to the ITU-T, a US position of support for HO and LO virtual concatenation.

References

- [1] T1X1.5/99-194, *Justification of Virtual Concatenation of SONET SPEs*, Lucent Technologies and Nortel Networks, T1X1.5 Meeting, Nashville, TN, July 26th-29th, 1999.
- [2] T1X1.5/99-XX, *Justification of Virtual Concatenation of SONET SPEs*, Nortel Networks, T1X1.5 Meeting, October 4th-8th, 1999.

Appendix A

The following is the text proposed at the recent ITU-T SG15 Q11 meeting for inclusion into G.707: -

11.2 Virtual concatenation of X VC-3/4s (VC-3/4-Xv, X = 1 ... 256)

A VC-3/4-Xv provides a contiguous payload area of X Container-3/4 (VC-3/4-Xc) with a payload capacity of $X \cdot 49536 / 149760$ kbit/s as shown in Figure 11-2 and 11-3. The container is mapped in X individual VC-3/4s which form the VC-3/4-Xv. Each VC-3/4 has its own POH as specified in subclause 9.3.1. The H4 POH byte is used for the virtual concatenation specific sequence and multi-frame indication as defined below.




FIGURE 11-2/G.707

VC-3-Xv structure




FIGURE 11-3/G.707

VC-4-Xv structure

Each VC-3/4 of the VC-3/4-Xv is transported individually through the network. Due to different propagation delay of the VC-3/4s a differential delay will occur between the individual VC-3/4s. This differential delay has to be compensated and the individual VC-3/4s have to be realigned for access to the contiguous payload area. The realignment process has to cover at least a differential delay of 125 μ s.

A two-stage 512 ms multi-frame is introduced to cover differential delays of 125 μ s and above (up to 256 ms). The first stage uses H4, bits 5-8 for the 4 bit multi-frame indicator (MFI1). MFI1 is incremented every basic frame and counts from 0 to 15. For the 8 bit multi-frame indicator of the second stage (MFI2), H4, bits 1-4 in frame 0 (MFI2 bits 1-4) and 1 (MFI2 bits 5-8) of the first multi-frame are used (see Table 9). MFI2 is incremented once every multi-frame of the first stage and counts from 0 to 255. The resulting overall multi-frame is 4096 frames (= 512 ms) long. The sequence indicator SQ identifies the sequence/order in which the individual VC-3/4s of the VC-3/4-Xv are combined to form the contiguous container VC-3/4-Xc as shown in Figure 11-4. Each VC-3/4 of a VC-3/4-Xv has a fixed unique sequence number in the range of 0 to (X-1). The VC-3/4 transporting the first time slot of the VC-3/4-Xc has the sequence number 0, the VC-3/4 transporting the second time slot the sequence number 1 and so on up to the VC-3/4 transporting time slot X of the VC-3/4-Xc with the sequence number (X-1). The sequence number is fixed assigned and not configurable. It allows one to check the correct constitution of the VC-3/4-Xv

without using the trace. The 8-bit sequence number (which supports values of X up to 256) is transported in bits 1 to 4 of the H4 bytes, using frame 14 (SQ bits 1-4) and 15 (SQ bits 5-8) of the first multi-frame stage as shown in Table 9.

TABLE 9/G.707

VC-3/4-Xv sequence and multi frame indicator H4 coding

H4 Byte								1 st multi-frame number	2 nd multi-frame number
Bit 1	Bit 2	Bit 3	Bit 4	Bit 5	Bit 6	Bit 7	Bit 8		
				1 st multi-frame indicator MFI1 (bits 1-4)					
Sequence indicator MSB (bits 1-4)				1	1	1	0	14	n-1
Sequence indicator LSB (bits 5-8)				1	1	1	1	15	
2 nd multi-frame indicator MFI2 MSB (bits 1-4)				0	0	0	0	0	n
2 nd multi-frame indicator MFI2 LSB (bits 5-8)				0	0	0	1	1	
Reserved ("0000")				0	0	1	0	2	
Reserved ("0000")				0	0	1	1	3	
Reserved ("0000")				0	1	0	0	4	
Reserved ("0000")				0	1	0	1	5	
Reserved ("0000")				0	1	1	0	6	
Reserved ("0000")				0	1	1	1	7	
Reserved ("0000")				1	0	0	0	8	
Reserved ("0000")				1	0	0	1	9	
Reserved ("0000")				1	0	1	0	10	
Reserved ("0000")				1	0	1	1	11	
Reserved ("0000")				1	1	0	0	12	
Reserved ("0000")				1	1	0	1	13	
Sequence indicator SQ MSB (bits 1-4)				1	1	1	0	14	
Sequence indicator SQ LSB (bits 5-8)				1	1	1	1	15	
2 nd multi-frame indicator MFI2 MSB (bits 1-4)				0	0	0	0	0	n+1
2 nd multi-frame indicator MFI2 LSB (bits 5-8)				0	0	0	1	1	
Reserved ("0000")				0	0	1	0	2	




FIGURE 11-4/G.707

VC-3/4-Xv multiframe and sequence indicator

11.3 Contiguous concatenation of X VC-2s in a higher order VC-3 (VC-2-Xc, X = 1 ... 7)

A VC-2-Xc provides a payload area of X Container-2 as shown in Figure 11-5. One common set of POH, corresponding to the POH of the first VC-2, is used for the whole VC-2-Xc (e.g. the BIP-2 covers all $428 \cdot X$ bytes of the VC-2-Xc). The POH positions corresponding to VC-2#2 to VC-2#X are fixed stuff.

The VC-2-Xc is located in X contiguous TU-2 in a higher order VC-3. The first column of the VC-2-Xc is always located in the first TU-2. The pointer of this first TU-2 indicates the position of the V5 POH byte of the VC-2-Xc. The pointers of the TU-2 #2 to #X are set to the

concatenation indication (see Figure 8-10) to indicate the contiguous concatenated payload. Pointer justification is performed in common for the X concatenated TU-2s and X stuffing bytes are used.

With allowed values of X between 1 and 7, the VC-2-Xc provides a payload capacity between 6 784 kbit/s and 47 488 kbit/s in steps of 6 784 kbit/s.




FIGURE 11-5/G.707

VC-2-Xc structure

11.4 Virtual concatenation of X VC-2s in a higher order VC-4 (VC-2-Xv, X = 1 ... 21)

A VC-2-Xv provides a payload area of X Container-2 as shown in Figure 11-6. The container is mapped in X individual VC-2s which form the VC-2-Xv. Each VC-2 has its own POH. The individual VC-2 of a VC-2-Xv are transported in higher order VC-4s through the network. Due to the individual transport the order and the alignment of the VC-2s will change. At the termination the individual VC-2s have to be rearranged and realigned in order to re-establish the contiguous concatenated container. Overhead supporting these functions is for further study. In order to minimize the differential delay between the individual VC-2s of the VC-2-Xv, the individual VC-2s shall be transported in the same higher order VC-4. The maximum differential delay is for further study.

With allowed values of X between 1 and 21, the VC-2-Xv provides a payload capacity between 6 784 kbit/s and 142 464 kbit/s in steps of 6 784 kbit/s.




FIGURE 11-6/G.707

VC-2-Xv structure

Annex A

(This Annex has been deleted)

Appendix B

Measurement of Differential Delay

In scenarios as outlined above the equipment which is receiving the virtual concatenated SPE/VCs must be capable of measuring the differential delay and determining which of the SPEs from a concatenated group is the first to arrive. It is also important to know when the maximum tolerable differential delay has been reached.

To determine the magnitude of differential delay requires assigning one of the SPEs as a datum and measuring the relative position of the others to it. This implies a plus and minus range. This will also allow determination of which SPE arrives first. Fig 2 illustrates two SPEs (A and B) offset in time by three frames (time runs from bottom to top of the diagram), A arriving before B. Note that to obtain consistent results over the code wrap around region the code must be interpreted as 2's Complement.

SPE #A			SPE/#B			A-B		
FrmNo	H4/b5-8		2s Comp	FrmNo	H4/b5-8		2s Comp	
15	1111	-1	12	1100	-4	3		
14	1110	-2	11	1011	-5	3		
13	1101	-3	10	1010	-6	3		
12	1100	-4	09	1001	-7	3		
11	1011	-5	08	1000	-8	3		
10	1010	-6	07	0111	7	3		
09	1001	-7	06	0110	6	3		
08	1000	-8	05	0101	5	3		
07	0111	7	04	0100	4	3		
06	0110	6	03	0011	3	3		
05	0101	5	02	0010	2	3		
04	0100	4	01	0001	1	3		
03	0011	3	00	0000	0	3		
02	0010	2	15	1111	-1	3		
01	0001	1	14	1110	-2	3		
00	0000	0	13	1101	-3	3		
15	1111	-1	12	1100	-4	3		
14	1110	-2	11	1011	-5	3		
13	1101	-3	10	1010	-6	3		
12	1100	-4	09	1001	-7	3		

Fig 2: Measurement of Differential Delay

With the a multi-frame indicator of four bits used to align SPEs from the same concatenated group a total delay of $\pm 1\text{ms}$ is possible, but beyond this it is only possible to detect misalignment by corrupted data.

Detection of Excessive Differential Delay

It is essential for the system to be able to detect when the differential delay exceeds what it can handle. This limit will correspond to the size of the buffers used in the re-alignment process (Upper Memory Limit (UML) and Lower Memory Limit (LML) in Fig 3 below).

For differential delays greater than that that can be handled by the buffers the system must detect this and raise an alarm and perform consequent action (to be determined). To do this means that the counter range must be greater than the buffer range. In addition the counter range must be sufficiently large to prevent aliasing.

Aliasing

Consider the binary sequence in Fig 3 below. The three bit binary count that starts at 000, increments up the page, and decrements down the page. Assume that this represents the range of the counter addressing the buffer (i.e. differential delay) and that the buffer has an Upper Memory Limit (UML) of 001 and a Lower Memory Limit (LML) of 110. If a differential delay greater than the buffer size occurs then the counter will be beyond the buffer limit (e.g. 011) and an alarm/consequent action will result. If a large enough differential delay occurs to make the count go to 110 then this code is within the buffer range (LML), no alarm/consequent action and data corruption occurs! The binary codes in **Bold** are legitimate memory ranges.

To minimise the problem of aliasing the counter range must be considerably greater than the memory address range.

010			
001			
000	Alias		
111	Range		
110	Same as LML		
101			
100	Detect		
011	Range		
010			
001	UML		
000	Start Working		
111	Range		
110	LML		
101			
100	Detect		
011	Range		
010			
001	Same as UML		
000	Alias		
111	Range		
110			
101			

Fig 3

- End of the document -